12th International Conference on Information Fusion
Seattle, WA, USA, July 6-9, 2009

Track-Before-Detect for Sensors with Complex Measurements

S. J. Davey† B. Cheung†⋆ M. G. Rutten†
†Defence Science and Technology Organisation, Australia
⋆School of Elec Engineering, The University of Adelaide, Australia
Email: {Samuel.Davey, Mark.Rutten, Brian.Cheung}@DSTO.Defence.gov.au

Abstract – *Track-Before-Detect (TkBD) is a paradigm that combines the target detection and estimation processes that are usually sequentially applied to sensor data in a conventional system. Under TkBD the single frame detector is removed and the tracker is supplied with the whole sensor image. Detection decisions are then shifted to the output of the tracker which is able to use temporal correlation to improve the decision performance. A fundamental measure used by most TkBD approaches is the likelihood ratio of the sensor data and this is formed as the product over individual cell likelihoods under the assumption of spatially independent noise. However, that approach exploits only the envelope of the known sensor point spread function. This article presents an approach for determining the data likelihood ratio that also includes phase information. This alternative likelihood ratio formulation is shown to both improve the discrimination of targets from noise and reduce the computation overhead of the algorithm.*

1 Introduction

Traditional tracking algorithms are designed assuming that the sensor provides a set of point measurements at each scan. The tracking algorithm links measurements across time and estimates parameters of interest. However, a practical sensor may provide a data image, where each pixel corresponds to the received power in a particular spatial location (e.g. range bins and azimuth beams). In this case, the common approach is to apply a threshold to the data and to treat those cells that exceed the threshold as point measurements, perhaps using interpolation methods to improve accuracy. This is acceptable if the signal to noise ratio (SNR) is high. For low SNR targets the threshold must be low to allow sufficient probability of target detection. A low threshold also gives a high rate of false detections, which causes the tracker to form false tracks. An alternative approach, referred to as track-before-detect (TkBD), is to supply the tracker with all of the sensor data without applying a threshold. This improves track accuracy and allows the tracker to follow low SNR targets.

The key difference between TkBD algorithms and conventional trackers is the measurement model. Conventional trackers use a point measurement model that typically assumes that the observations are kinematic states of the target corrupted with additive (usually Gaussian) noise (e.g. [1, 3]). However, a TkBD algorithm requires a statistical model of the signal received in each sensor pixel.

Earlier TkBD efforts focussed on optical sensors [2, 6, 7, 13] and assumed that the sensor noise was Gaussian. More recent efforts have considered the radar problem where the pixel intensity is either Ricean or Rayleigh distributed depending on the presence or absence of targets [4, 9, 10, 11]. The overall likelihood ratio is formed by taking a product of the pixel likelihood ratios for all pixels. This product is a result of assuming that the noise is spatially uncorrelated. The Ricean and Rayleigh distributions are functions of only the envelope of the sensor data, so this approach does not use phase information.

This article demonstrates an alternative likelihood formulation that allows for spatially correlated noise and includes the phase information. The performance of a grid based recursive posterior filter [5, 14] is analysed for the traditional likelihood function using only envelope data and the alternative likelihood function that includes phase. The inclusion of phase information is shown to improve the detection and estimation performance of the filter. Furthermore, the complex likelihood function is an order of magnitude more efficient to evaluate than the envelope one.

2 Problem Definition

As in [9, Ch. 11], consider a sensor that collects a sequence of two-dimensional images (frames). When present, a target moves in the plane according to a known statistical process. Each pixel in a frame is a complex valued random variable, and the target contribution is via a known complex point spread function with an arbitrary and unknown phase offset.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
JUL 2009		**06-07-2009 to 09-07-2009**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Track-Before-Detect for Sensors with Complex Measurements	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Defence Science and Technology Organization (DSTO),School of Electrical Engineering,The University of Adelaide, Australia, ,**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002299. Presented at the International Conference on Information Fusion (12th) (Fusion 2009). Held in Seattle, Washington, on 6-9 July 2009. U.S. Government or Federal Rights License.

14. ABSTRACT
Track-Before-Detect (TkBD) is a paradigm that combines the target detection and estimation processes that are usually sequentially applied to sensor data in a conventional system. Under TkBD the single frame detector is removed and the tracker is supplied with the whole sensor image. Detection decisions are then shifted to the output of the tracker which is able to use temporal correlation to improve the decision performance. A fundamental measure used by most TkBD approaches is the likelihood ratio of the sensor data and this is formed as the product over individual cell likelihoods under the assumption of spatially independent noise. However, that approach exploits only the envelope of the known sensor point spread function. This article presents an approach for determining the data likelihood ratio that also includes phase information. This alternative likelihood ratio formulation is shown to both improve the discrimination of targets from noise and reduce the computation overhead of the algorithm.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Public Release**	**8**	

2.1 Target Model

For simplicity of notation, assume a discrete time model, with a fixed frame period, T. The target state at frame k, $\mathbf{x}_k$, consists of position and velocity in the plane, and the intensity of the returned signal, I_k, i.e.

$$\mathbf{x}_k = \begin{bmatrix} x_k & \dot{x}_k & y_k & \dot{y}_k & I_k \end{bmatrix}^{\mathsf{T}}. \quad (1)$$

The evolution of the state is modelled by the linear stochastic process

$$\mathbf{x}_k = \mathsf{F}\mathbf{x}_{k-1} + \mathbf{v}_k, \quad (2)$$

where $\mathbf{v}_k$ is a noise process and the transition matrix is given by

$$\mathsf{F} = \begin{bmatrix} \mathsf{F}_s & 0 & 0 \\ 0 & \mathsf{F}_s & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathsf{F}_s = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}. \quad (3)$$

The process noise, $\mathbf{v}_k$, is the usual Gaussian random variable with covariance Q given by

$$\mathsf{Q} = \begin{bmatrix} \mathsf{Q}_s & 0 & 0 \\ 0 & \mathsf{Q}_s & 0 \\ 0 & 0 & q_i T \end{bmatrix}, \quad \mathsf{Q}_s = q_s \begin{bmatrix} T^3/3 & T^2/2 \\ T^2/2 & T \end{bmatrix}, \quad (4)$$

where q_s is the power spectral density of the acceleration noise in the spatial dimensions and q_i is the power spectral density of the noise in the rate of change of target return intensity.

2.2 Measurement Model

The measurement at each time is an image of an arbitrary dimensionality with N pixels indexed by $i = 1 \ldots N$. Denote the complex response in the ith pixel of the kth frame as z_k^i and let $\boldsymbol{z}_k$ be a stacked vector of all the pixel responses of the kth frame. The target signal at the ith pixel is denoted by $h^i(\mathbf{x}_k)$ and $\mathbf{h}(\mathbf{x}_k)$ is a stacked vector of the target signal over all of the pixels. This target signal may be the point spread function (psf) of the sensor, the target signature (range profile) or a combination of both. It is shifted to be centred over the target state and scaled by the target amplitude and assumed to be known.

Assuming additive noise, the measurement equation is

$$\boldsymbol{z}_k = \exp\{j\phi\}\mathbf{h}(\mathbf{x}_k) + \mathbf{n}_k, \quad (5)$$

where ϕ is an unknown phase shift, uniformly distributed over $[0, 2\pi]$, and $\mathbf{n}_k$ is a stacked vector of the (complex) pixel noise signals. Assume that $\mathbf{n}_k$ is complex Gaussian with a known covariance matrix, R. For this paper, the noise will be assumed to be spatially independent, so R is an identity matrix scaled by the known noise variance, σ^2.

The target peak SNR quantifies the height of the peak of the target point spread function relative to the noise floor, and represents a measure of how easy it is to detect the target. The point spread function is assumed to be normalised such that the contribution to cell i is I_k when the target is located exactly on the sample point for the cell. Thus the peak SNR in dB is given by $20 \log \left\{ I_k / \sigma^2 \right\}$.

2.3 Envelope Likelihood

The standard treatment of this measurement model is to assume that the noise is spatially uncorrelated and to form the likelihood of its envelope (i.e. the magnitude of the data) [11, 4, 9]. Thus R is assumed to be an identity matrix scaled by σ^2.

The envelope likelihood considers only the magnitude of the sensor response at each pixel. The pixel magnitude will be Ricean distributed, if there is a target present, or Rayleigh distributed if there is no target [12]. Thus the pdf of $\left|z_k^i\right|$ is

$$p\left(\left|z_k^i\right| \middle| \mathbf{x}_k\right) = \frac{\left|z_k^i\right|}{\sigma^2} \exp\left(-\frac{\left|z_k^i\right|^2 + \left|h^i(\mathbf{x}_k)\right|^2}{2\sigma^2}\right) \times \mathsf{I}_0\left(\frac{\left|h^i(\mathbf{x}_k)\ z_k^i\right|}{\sigma^2}\right), \quad (6)$$

if the target is present, or

$$p\left(\left|z_k^i\right|\right) = \frac{\left|z_k^i\right|}{\sigma^2} \exp\left(-\frac{\left|z_k^i\right|^2}{2\sigma^2}\right), \quad (7)$$

if there is no target, where $\mathsf{I}_0(\cdot)$ is the modified Bessel function.

Note that the measurement probability implied by (5) depends on the phase shift, ϕ. The pdf (6) is the marginal distribution achieved by the integral

$$p\left(\left|z_k^i\right| \middle| \mathbf{x}_k\right) = \int_0^{2\pi} p\left(\left|z_k^i\right| \middle| \mathbf{x}_k, \phi\right) p(\phi) d\phi. \quad (8)$$

The Bessel function arises as a result of this integral [8].

The likelihood ratio for the magnitude response at pixel i is thus

$$\mathcal{L}\left(\left|z_k^i\right| \middle| \mathbf{x}_k\right) \equiv \frac{p\left(\left|z_k^i\right| \middle| \mathbf{x}_k\right)}{p\left(\left|z_k^i\right|\right)} = \exp\left(\frac{-\left|h^i(\mathbf{x}_k)\right|^2}{2\sigma^2}\right) \mathsf{I}_0\left(\frac{\left|h^i(\mathbf{x}_k)\ z_k^i\right|}{\sigma^2}\right). \quad (9)$$

Since the pixels are assumed to be conditionally independent, the likelihood of the envelope of the whole image is simply the product over the pixels

$$\mathcal{L}_e\left(|z_k| \middle| \mathbf{x}_k\right) = \prod_{i=1}^{N} \mathcal{L}\left(\left|z_k^i\right| \middle| \mathbf{x}_k\right), = \exp\left(\frac{-\mathbf{h}(\mathbf{x}_k)^{\mathsf{H}}\mathbf{h}(\mathbf{x}_k)}{2\sigma^2}\right) \prod_{i=1}^{N} \mathsf{I}_0\left(\frac{\left|h^i(\mathbf{x}_k)\ z_k^i\right|}{\sigma^2}\right), \quad (10)$$

where $^{\mathsf{H}}$ denotes the Hermition conjugate.

Typically this product is truncated to only include those pixels that have a significant target contribution, $h^i(\mathbf{x}_k)$, to avoid unnecessary computation [11].

If a prior distribution is assumed for the target intensity, $p(I_k)$, then an intensity independent marginal likelihood is given by

$$\bar{\mathcal{L}}_e\left(\boldsymbol{z}_k|\mathbf{x}_k\right) = \int_0^\infty \mathcal{L}_e\left(\boldsymbol{z}_k|\mathbf{x}_k\right) p(I_k) dI_k. \qquad (11)$$

This integral may be numerically approximated using a summation.

An important implication of the above method is that marginalisation over the random phase shift, ϕ, has been repeated independently for each pixel. This is equivalent to assuming that there is a different phase shift for every pixel, ϕ_k^i. This is not the case, and so the envelope likelihood has not captured all of the available information.

3 Complex Likelihood

The factorised form of the envelope likelihood results in a loss of information because it fails to include the phase of the sensor data. It also requires the calculation of a Bessel function for each pixel with a significant psf contribution. Experience shows that the evaluation of these Bessel functions is by far the most computationally expensive part of TkBD algorithms that use this likelihood.

Consider now the joint likelihood of the whole sensor image. For compactness, the notation will omit the dependence of the psf on the target state, omit the frame index k, and let $s = \exp\{j\phi\}$ with $s^\star$ its conjugate. This likelihood defines a matched filter which can account for both spatially correlated noise and the phase response of the psf. The derivation of this likelihood is now presented.

The probability density of $\boldsymbol{z}$ when a target is present is

$$P(\boldsymbol{z}|\text{target}, \phi) = \frac{1}{|2\pi\mathsf{R}|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2}\left(\boldsymbol{z} - s\mathbf{h}\right)^{\mathsf{H}} \mathsf{R}^{-1}\left(\boldsymbol{z} - s\mathbf{h}\right)\right\}, \quad (12)$$

and when a target is absent is

$$P(\boldsymbol{z}|\text{no target}) = \frac{1}{|2\pi\mathsf{R}|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2}\boldsymbol{z}^{\mathsf{H}}\mathsf{R}^{-1}\boldsymbol{z}\right\}, \quad (13)$$

since $\mathbf{n}$ is Gaussian with covariance R.

Thus the likelihood ratio is

$$\begin{aligned}
&\mathcal{L}(\boldsymbol{z}|\boldsymbol{x}, \phi) \\
&= \exp\left\{-\frac{1}{2}\left(\boldsymbol{z} - s\mathbf{h}\right)^{\mathsf{H}} \mathsf{R}^{-1}\left(\boldsymbol{z} - s\mathbf{h}\right) + \frac{1}{2}\boldsymbol{z}^{\mathsf{H}}\mathsf{R}^{-1}\boldsymbol{z}\right\}, \\
&= \exp\left\{\frac{1}{2}\boldsymbol{z}^{\mathsf{H}}\mathsf{R}^{-1}s\mathbf{h} + \frac{1}{2}s^\star\mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}\boldsymbol{z} - \frac{1}{2}s^\star\mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}s\mathbf{h}\right\}, \\
&= \Lambda \exp\left\{\frac{1}{2}s\boldsymbol{z}^{\mathsf{H}}\mathsf{R}^{-1}\mathbf{h} + \frac{1}{2}s^\star\mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}\boldsymbol{z}\right\}, \\
&= \Lambda \exp\left\{\frac{1}{2}s\xi^\star + \frac{1}{2}s^\star\xi\right\}, \qquad (14)
\end{aligned}$$

where

$$\begin{aligned}
\Lambda &\equiv \exp\left\{-\frac{1}{2}\mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}\mathbf{h}\right\}, \\
\xi &\equiv \mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}\boldsymbol{z} \equiv \Xi\exp\{j\theta\}.
\end{aligned}$$

Algebraic manipulation yields

$$\begin{aligned}
&\mathcal{L}(\boldsymbol{z}|\boldsymbol{x}, \phi) \\
&= \Lambda \exp\left\{\frac{1}{2}(\cos\phi + j\sin\phi)\xi^\star + \frac{1}{2}(\cos\phi - j\sin\phi)\xi\right\}, \\
&= \Lambda \exp\left\{\frac{1}{2}\cos\phi[\xi^\star + \xi] + \frac{1}{2}j\sin\phi[\xi^\star - \xi]\right\}, \\
&= \Lambda \exp\left\{\Xi\cos\phi\cos\theta + \Xi\sin\phi\sin\theta\right\}, \\
&= \Lambda \exp\left\{\Xi\cos(\phi - \theta)\right\}. \qquad (15)
\end{aligned}$$

Marginalise over ϕ assuming a uniform distribution:

$$\begin{aligned}
\mathcal{L}(\boldsymbol{z}|\boldsymbol{x}) &= \int_0^{2\pi} \mathcal{L}(\boldsymbol{z}|\boldsymbol{x}, \phi) p(\phi) d\phi \\
&= \frac{\Lambda}{2\pi}\int_0^{2\pi} \exp\left\{\Xi\cos(\phi - \theta)\right\} d\phi \\
&= \exp\left\{-\frac{1}{2}\mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}\mathbf{h}\right\} \mathsf{I}_0\left(\left|\mathbf{h}^{\mathsf{H}}\mathsf{R}^{-1}\boldsymbol{z}\right|\right) \qquad (16)
\end{aligned}$$

If the noise is spatially uncorrelated with variance σ^2, then R is the identity matrix scaled by σ^2 and the likelihood simplifies to

$$\mathcal{L}(\boldsymbol{z}|\boldsymbol{x}) = \exp\left\{-\frac{\mathbf{h}^{\mathsf{H}}\mathbf{h}}{2\sigma^2}\right\} \mathsf{I}_0\left(\frac{\left|\mathbf{h}^{\mathsf{H}}\boldsymbol{z}\right|}{\sigma^2}\right) \qquad (17)$$

As with the envelope likelihood, the above expression is implicitly dependent on the target amplitude. One way to remove this dependence is to marginalise over it as in (11).

The likelihood expression is very similar to the envelope likelihood except that whereas the envelope likelihood includes a product of Bessel functions, (17) has a single Bessel function evaluated at the output of a matched filter. Because evaluating the Bessel functions is the main computation cost of many TkBD algorithms, one would expect a significant speed improvement by using (17) instead of the envelope likelihood. The key difference is that the matched filter form includes the phase profile of the psf which is lost in the envelope likelihood. This extra information should lead to superior performance.

4 State Posterior Recursion

The test TkBD algorithm is a discretised numerical approximation to the exact Bayesian posterior recursion. The algorithm uses a fixed grid in the state space and propagates the posterior probability mass of the target state over this grid. The mode of the posterior is used as the estimate in

this paper, since it is likely that the distribution will be multimodal, and the mean may be sensitive to distant low likelihood modes.

The posterior pdf of the target state can be recursively determined using the well known Bayesian relationship

$$p(\mathbf{x}_k|Z_k) \propto p(z_k|\mathbf{x}_k) \int p(\mathbf{x}_k|\mathbf{x}_{k-1}) p(\mathbf{x}_{k-1}|Z_{k-1}) d\mathbf{x}_{k-1}, \tag{18}$$

where $Z_k \equiv \boldsymbol{z}_1 \ldots \boldsymbol{z}_k$.

The estimator in this article uses a direct approximation to (18) based on a discretisation of the state space. Choose a uniformly spaced set of states, $\mathcal{X}$ (which is not necessarily related to the discrete measurement function) such that

$$\mathbf{x}_k(q,r,s,t) = \begin{bmatrix} \Delta_x q & \frac{\Delta_x}{T} r & \Delta_y s & \frac{\Delta_y}{T} t & I_k \end{bmatrix}^T, \tag{19}$$

for some integers q,r,s and t.

(18) can then be approximated by

$$p(\mathbf{x}_k|Z_k) \approx K \bar{\mathcal{L}}(z_k|\mathbf{x}_k) \sum_{\mathbf{x}_{k-1} \in \mathcal{X}} p(\mathbf{x}_k|\mathbf{x}_{k-1}) p(\mathbf{x}_{k-1}|Z_{k-1}), \tag{20}$$

where K is a normalising constant and

$$\sum_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}) \equiv \sum_{q=-\infty}^{\infty} \sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} F\Big(\mathbf{x}(q,r,s,t)\Big). \tag{21}$$

The approximation is exact in the limit as $\mathcal{X}$ approaches $\Re^4$ for non-negative continuous functions, as used here.

The first term in (20) is the intensity independent marginal likelihood, defined by (11) and using either the envelope likelihood or the matched likelihood.

The state evolution term in (20), $p(\mathbf{x}_k|\mathbf{x}_{k-1})$, is a four dimensional matrix. The dynamics determines the centre of the distribution, i.e. $\mathsf{F}\mathbf{x}_{k-1}$, and the shape is prescribed by the probability mass of the process noise, $\mathbf{v}_k$, which is now discrete on $\mathcal{X}$. In order to minimise the computation burden of (20), the probability mass of $\mathbf{v}_k$ is assumed to be zero outside a compact region centred on the origin.

The discrete state space is augmented with a null state, $\emptyset$, to indicate the possibility that there is no target. Denote the probability of target *death* as P_d, and the probability of target *birth* as P_b. Then the evolution probability in (20) is given by

$$p(\mathbf{x}_k|\mathbf{x}_{k-1}) = \begin{cases} 1 - P_b & \mathbf{x}_k = \emptyset, \mathbf{x}_{k-1} = \emptyset, \\ P_d & \mathbf{x}_k = \emptyset, \mathbf{x}_{k-1} \neq \emptyset, \\ P_b/|\mathcal{X}| & \mathbf{x}_k \neq \emptyset, \mathbf{x}_{k-1} = \emptyset, \\ (1 - P_d) p_v(\tilde{v}_k) & \mathbf{x}_k \neq \emptyset, \mathbf{x}_{k-1} \neq \emptyset, \end{cases} \tag{22}$$

where $|\mathcal{X}|$ is the number of discrete states in $\mathcal{X}$, $\tilde{v}_k = \mathbf{x}_k - F\mathbf{x}_{k-1}$ and $p_v(\tilde{v}_k)$ is the process noise probability distribution (refer to section 2.1).

The parameters P_b and P_d are tuning parameters and may be adjusted to optimise detection performance. The selection of the state space, $\mathcal{X}$, is a trade-off between estimation accuracy, which improves with finer resolution, and computation requirement, which increases with $|\mathcal{X}|$. The process noise pmf also affects estimation accuracy, as well as providing some capacity to handle model mismatch between the assumed target model and the true target motion. The algorithm is initialised with $p(\mathbf{x}_0 = \emptyset) = 1$ and $p(\mathbf{x}_0) = 0 \ \forall \, \mathbf{x}_0 \neq \emptyset$.

Once the pdf of the state has been evaluated, a state estimate can be obtained by selecting the state with the highest probability. In the event that this state is the null state, then the algorithm reports that there is no target. To account for the case where the pdf has a peak that is spread over several grid cells, the implementation used in this article finds the highest probability non-null state and accumulates the probability in the adjacent cells. If the accumulated probability is higher than the null state probability, then a detection is reported.

In the experimental analysis, it will be seen that implementations of the Bayesian posterior recursion using the two likelihood alternatives have substantially different computation costs. In order to remove the effects of implementation efficiency, the computation cost of the algorithm is now considered analytically.

Assume that the state space $\mathcal{X}$ is a regular grid such that $|\mathcal{X}| = A_X A_{\dot{X}}$, where A_X is the number of position states in $\mathcal{X}$ and $A_{\dot{X}}$ is the number of velocity states. Let V denote the number of elements of $\mathcal{X}$ for which the process noise has a non-zero probability mass.

The likelihood is independent of the target speed and so an efficient implementation requires A_X likelihood calculations, each of which has a cost $C_{\mathcal{L}}$.

For each state locale, (20) requires V products and accumulates under the sum with a further two multiplies outside the sum. Determining the scaling constant requires a sum over the whole state space. Let $C_\times$ and C_+ denote the cost of a multiplication and a sum respectively.

Thus the total cost of a single time slice of the estimator is

$$\begin{aligned} C &= A_X C_{\mathcal{L}} + A_X A_{\dot{X}} C_+ \\ &\quad + A_X A_{\dot{X}} \Big\{ (V+2) C_\times + V C_+ \Big\}, \\ &\approx A_X C_{\mathcal{L}} + A_X A_{\dot{X}} V \left(C_\times + C_+ \right), \end{aligned} \tag{23}$$

since $V \gg 1$.

If the state space has a fixed physical extent in space and velocity and a symmetrical sampling spatial frequency f_s then the cost can be written as

$$C \approx (f_s)^2 a_1 C_{\mathcal{L}} + (f_s)^8 a_2 \left(C_\times + C_+ \right), \tag{24}$$

where a_1 and a_2 are constants whose value are a function of the physical extent of the state space in position and velocity.

So the algorithm cost contains a term that has $O(N^2)$ complexity in the grid resolution due to the likelihoods and a term that has $O(N^8)$ complexity due to the dynamic model.

It is intuitive that the cost of calculating the envelope likelihood is much higher than that of the matched likelihood. Although the second term is high in complexity, it is

amenable to pipeline processing and may be relatively cheap for problems over a small spatial area.

5 Performance Evaluation

Bayesian algorithms constructed using the two likelihood functions are now compared via Monte Carlo simulation. The test scenarios consist of a small surveillance volume containing either a single straight line target, or no target. Although the straight line target assumption is rather limiting, this paper is focussed on the measurement model, not the target dynamics, so there is no requirement for more challenging targets.

A TkBD algorithm is primarily a detector that exploits a target dynamics model to achieve integration gain. Therefore the most important measure of performance is how well the algorithm discriminates between the presence and the absence of a target. Of secondary interest is the processing resource required to employ the algorithm. This quantity is highly dependent on the implementation. However, the variation in implementation cost should reflect the complexity model described above. The estimation error is also of interest, although it is much less important than the detection performance and processing cost.

5.1 Simulation Details

The simulation scenario uses a 30×45 pixel scan with unit variance complex normal noise. When present, the target signature is the spectral response of a Hann window centred at the target location. The target moves with a fixed speed of 1 pixels per frame. For each Monte Carlo trial a heading is uniformly sampled between 0 degrees (East) and 45 degrees (North-East).

The target state space was defined as the measurement domain in the spatial dimensions and ±3 pixels per frame in each velocity dimension. The resulting 4D state space has a size of $30 \times 7 \times 45 \times 7$.

Scalloping loss may reduce the sensitivity of the algorithm to a target that is located part way between sample points, resulting in a potential variation in performance depending on the relative position on the target with respect to the grid. In order to average over this potential variation, the initial target position for each Monte Carlo trial was randomly sampled from the range $2.5 \ldots 3$ independently in X and Y.

The Bayesian system model comprises a target dynamics model, a measurement model, and a target birth-death model. Of these, the measurement model is assumed to be known, and the target model is assumed to be fixed. The target model is specified by the process noise, which is assumed zero outside of a $3\times3\times3\times3$ region with probability 0.7 at the origin and the remaining 0.3 spread equally over the other 80 elements.

The parameters of the birth-death model are the probability of new target birth and the probability of old target death. For this comparison, it was found that using a very small probability of death gave best performance for all birth probabilities, so the probability of death was fixed at $P_d = 10^{-5}$.

The probability of birth, P_b, was varied to trace out a Receiver Operating Characteristic (ROC) curve. For each value of P_b, 100 Monte Carlo trials were performed with a target present and 100 without a target.

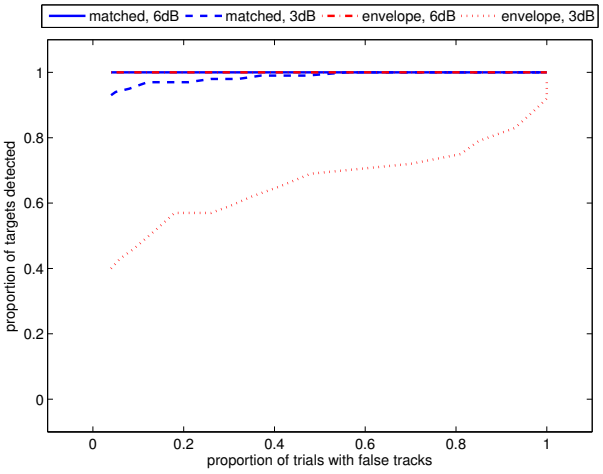

Figure 1: Overall ROC

5.2 Performance Measures and Results

Five Measures Of Performance (MOPs) were defined for the comparison: overall ROC; per-scan ROC; per-scan detection as a function of SNR; RMS position error and computation resource.

For TkBD, the main focus is on detection of targets, thus the first three measures are the most important. Of lesser interest is the estimation accuracy and computation cost, provided that the latter is not unreasonably high.

When a target scenario was tested, the target was declared detected on a particular scan if the TkBD algorithm produced an output track for that scan that was sufficiently close to the true target location. The association gate used only the target position and had a radius of 2 pixels.

5.2.1 Overall ROC

The proportion of targets detected is defined as the fraction of Monte Carlo trials where the TkBD algorithm produced a track that was within the association gate of the target location for at least one scan. The proportion of trials with false tracks is defined as the fraction of Monte Carlo trials where the TkBD algorithm produced at least one false track. This statistic only includes trials where there was no target present.

The overall ROC was generated by plotting the proportion of targets detected as a function of the proportion of trials with false tracks. This MOP quantifies how well the different likelihood ratios discriminate between data that contains a target and data that does not.

Figure 1 shows the overall ROC curves for both of the likelihood functions at target SNR values of 6dB and 3dB.

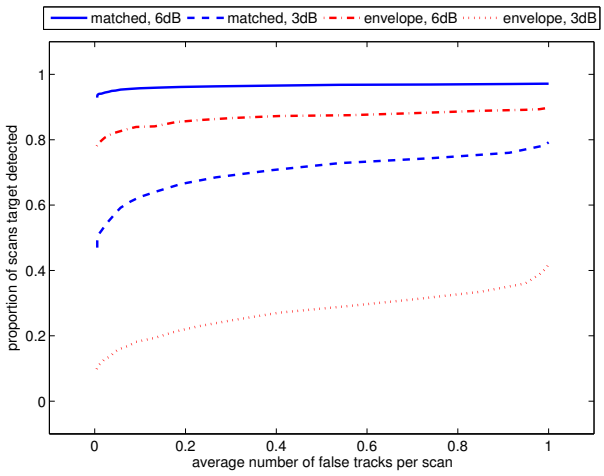

Figure 2: Per-scan ROC

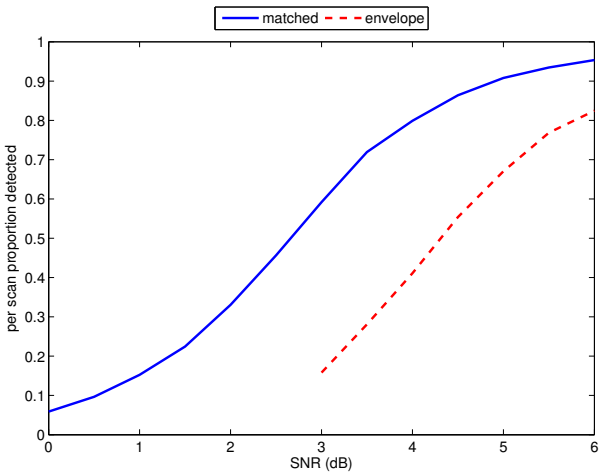

Figure 3: Per-scan detection as a function of SNR

Both likelihood functions detect all targets at 6dB, even for very low false alarm rates. At 3dB, the matched likelihood still provides very good detection performance, detecting approximately 95% of targets at very low false alarm. However, the envelope likelihood is unreliable at 3dB and only detects around half of the targets. The matched likelihood clearly provides superior overall detection performance.

5.2.2 Per-scan ROC

The overall ROC does not include any information about how long the algorithm takes to detect a target, or how long false tracks persist. In order to limit the number of MOPs these considerations are summarised by empirically estimating the per-scan probability of detection and false alarm.

The per-scan proportion of targets detected is defined as the fraction of scans where the TkBD algorithm produced a track that was within the association gate of the target location. An algorithm that quickly initiates a track and maintains the track will have a high per-scan proportion detected, whereas one that has a long initiation delay or quickly diverges will have a low one.

The per-scan false track rate is defined as the average number of false tracks reported by the TkBD algorithm per scan. This statistic summarises both the number and length of false tracks and is simply the ratio of the total number of false reports from the TkBD output to the total number of simulated frames.

The per-scan ROC was generated by plotting the per-scan proportion detected as a function of the per-scan false track rate and is shown in figure 2.

At 6dB, the matched likelihood per-scan proportion detected is close to 0.95, reflecting that the algorithm establishes a track (on average) by the second scan. The envelope likelihood has a slightly lower per-scan proportion detected, reflecting a track establishment delay of around 2-3 scans at 6dB.

At 3dB, the matched likelihood per-scan proportion detected is significantly lower than the overall proportion detected. This indicates that the algorithm now takes a more substantial time to establish track, on average 8 scans at low false alarm rates. The envelope likelihood per-scan proportion detected is less than half the overall proportion detected. This is because many of the targets it detects produce only short track segments.

As before, the matched likelihood is clearly superior.

5.2.3 Per-scan detection as a function of SNR

The ROC curves shown in figures 1 and 2 correspond to two SNR values. For this MOP, the per-scan false track rate was held fixed and the per-scan proportion of targets detected was observed as a function of SNR. For this MOP the birth probability was fixed to achieve 0.05 false reports per scan.

The per-scan proportion detected is shown in figure 3. It demonstrates that the improved detection ability of the matched likelihood function is approximately equivalent to a gain of 1.5 dB in signal power.

5.2.4 RMS position error

The two dimensional position error was averaged over those frames when the target was detected. Figure 4 shows the RMS error for each of the likelihoods as a function of SNR. The probability of birth was again fixed to achieve 0.05 false reports per scan. The RMS error for the envelope likelihood was only measured at 3 dB and above since there are too few tracks to get reliable estimates below this.

As should be expected, the estimation error is lower when the target SNR is higher. The RMS error of the matched likelihood tracks is roughly equivalent to that of the envelope likelihood at 2 dB higher SNR.

5.2.5 Computation resource

Table 1 gives the computation resource spent by the two alternatives both in absolute CPU seconds and in relative terms. Profiling the algorithm code using the envelope likelihood function showed that more than 90 percent of the cpu time was spent in the likelihood calculation, at around 55

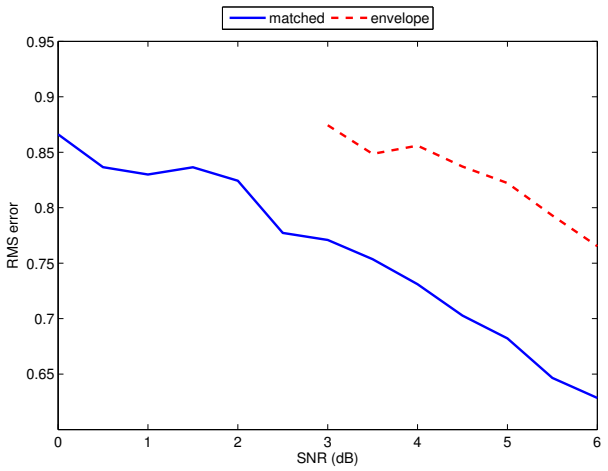

Figure 4: RMS error (pixels) as a function of SNR

likelihood	cpu time (s)	cpu time (ratio)
envelope	39,100	11
matched	3,570	1

Table 1: Computation requirements

seconds per monte carlo trial. In contrast, the calculations for the matched likelihood took only 1.4 seconds per trial, on average. This reduction is partly due to the number of Bessel function evaluations, which dropped by a factor of approximately 20. However, additional gain was made due to simplification of the inner loop in the Matlab code. This additional gain is not likely to be reflected in an implementation of the algorithm using a lower level software language. The nett result is approximately an order in magnitude speed improvement in the total algorithm.

The difference in computation resource required for an operational implementation of these algorithms would be highly dependent on how the Bessel functions were evaluated. If low-accuracy approximations to the likelihood were found to be sufficient, then the speed difference may be small. However, if higher accuracy Bessel evaluations were required, the order of magnitude speed difference may be realistic, since Matlab uses external optimised code to perform this task.

6 Summary

This paper demonstrated that the performance of a Bayesian Track-Before-Detect algorithm was significantly improved by making use of the phase signature of the target in addition to the amplitude signature. Intuitively, this phase information may be incorporated by using a matched filter tuned to the known complete point spread function of the sensor. This matched filter is in addition to those that may be used for range gating or other pre-tracker processing.

The use of phase information was found to be roughly equivalent to a 1.5 dB improvement in signal power, compared with the phase ignorant algorithm. The phase information not only improved detection and estimation performance, but also lead to a more efficient algorithm with an order of magnitude lower computation resource requirements.

References

[1] Y Bar-Shalom and X R Li. *Multitarget-multisensor tracking: principles and techniques*. YBS, Storrs, Conneticut, USA, 1995.

[2] Y. Barniv. Dynamic programming solution for detecting dim moving targets. *IEEE Transactions on Aerospace and Electronic Systems*, 21(1):144–156, January 1985.

[3] S. Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, Norwood, MA, 1999.

[4] Y. Boers and J. N. Driessen. Particle filter based detection for tracking. In *Proceedings of the American Control Conference*, pages 4393–4397, Arlington, VA, USA, June 2001.

[5] M. G. S. Bruno and J. M. F. Moura. Multiframe detector/tracker: Optimal performance. *IEEE Transactions on Aerospace and Electronic Systems*, 37(3):925–945, July 2001.

[6] R. M. Burczewski and N. C. Mohanty. Detection of moving optical objects. In *International Telemetering Conference*, pages 325–330, Los Angeles, California, 1978.

[7] P. S. Maybeck and D. E. Mercier. A target tracker using spatially distributed infrared measurements. *IEEE Transactions on Automatic Control*, 25(2):222–225, April 1980.

[8] H. V. Poor. *An Introduction to Signal Detection and Estimation*. Springer-Verlag, 1988.

[9] B. Ristic, S. Arulampalam, and N. J. Gordon. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004.

[10] M. G. Rutten, N. J. Gordon, and S. Maskell. Recursive track-before-detect with target amplitude fluctuations. *IEE Proceedings on Radar, Sonar and Navigation*, 152(5):345–352, October 2005.

[11] D. J. Salmond and H. Birch. A particle filter for track-before-detect. In *Proceedings of the American Control Conference*, pages 3755–3760, Arlington, VA, USA, June 2001.

[12] M. I. Skolnik. *Introduction to Radar Systems*. McGraw-Hill, third edition, 2001.

[13] M. C. Smith and E. M. Winter. On the Detection of Target Trajectories in a Mutli Target Environment. In *IEEE Conference on Decision and Control*, pages 1189–1194, San Diego, California, 1978.

[14] L. D. Stone, C. A. Barlow, and T. L. Corwin. *Bayesian Multiple Target Tracking*. Artech House, 1999.